

Secure or Insure?

Analyzing Network Security Games with Externalities

Nicolas Christin

Carnegie Mellon University, CyLab & INI

Collaborators

John Chuang (UC Berkeley),
Jens Grossklags (Penn State),
Benjamin Johnson (Carnegie Mellon)

Relevant Publications

- J. Grossklags, N. Christin, and J. Chuang. Secure or Insure? A Game-Theoretic Analysis of Information Security Games. [WWW'08](#).
- J. Grossklags, N. Christin, and J. Chuang. Security and Insurance Management in Networks with Heterogeneous Agents. [ACM EC'08](#).
- J. Grossklags, B. Johnson and N. Christin. When Information Improves Information Security. [FC'10](#).
- B. Johnson, J. Grossklags, N. Christin and J. Chuang. Are security experts useful? Bayesian Nash Equilibria for Network Security Games with Limited Information. [ESORICS'10](#).

Problem statement

Why do people and corporations not invest more in security?

- Users claim they have an interest in secure practices
- Security technology is (by and large) available for cheap
 - PGP, SSL, AES...
- Financial losses can be very large
 - Identity theft, loss of reputation

Thesis

- Economics can help understand and change user behavior
 - Everybody is on a network
 - Competitive environment
 - E.g., different ISPs, content providers, different divisions...
 - Strong externalities
 - Each user's security affects the whole network
 - Who should pay?
 - Development of criminal markets
 - Criminals very rational (see Willie Sutton): in it for the money
- Users not perfectly rational, but not random either
- Complement to technological solutions

Modeling methodology



1. Formal analysis

Today

- Game-theoretic predictions, selfishness vs. altruism
- Impact of various parameters

2. Experimental research

- Controlled lab and online experiments
- Behavioral modeling

3. Field data measurement

- Acquisition of attacker data (criminal markets goods, advertisements, ...)
- Acquisition of investment patterns

4. Testing intervention mechanisms

- Incentives, legal...

Formal analysis: Approach and contribution

- Variety of security threats and responses
 - Model most security interactions met in practice
 - Finite number of canonical security games
- Decouple security strategies
 - Self-protection investments (e.g., setting up a firewall)
 - Self-insurance coverage (e.g., archiving data as back up)
- Consider **network externalities**
 - Choice of strategy by a network participant affects other participants

Foundations of our security model

- Common and conflicting interests
 - Pure conflict, in which the interests of the two antagonists are completely opposed, is a special case. (Schelling, 1965)
- Public goods literature (Hirshleifer, 1983)
 - Non-rivalry: Simultaneous access
 - Non-excludability: Non-contributors can benefit
- Models adapted to reliability/security context (Varian, 2000)
 - Security can be interpreted as a public good
 - Differences to economics literature:

“Considerations of costs, benefits, and probability of failure become paramount, with income effects being a secondary concern”

Security actions

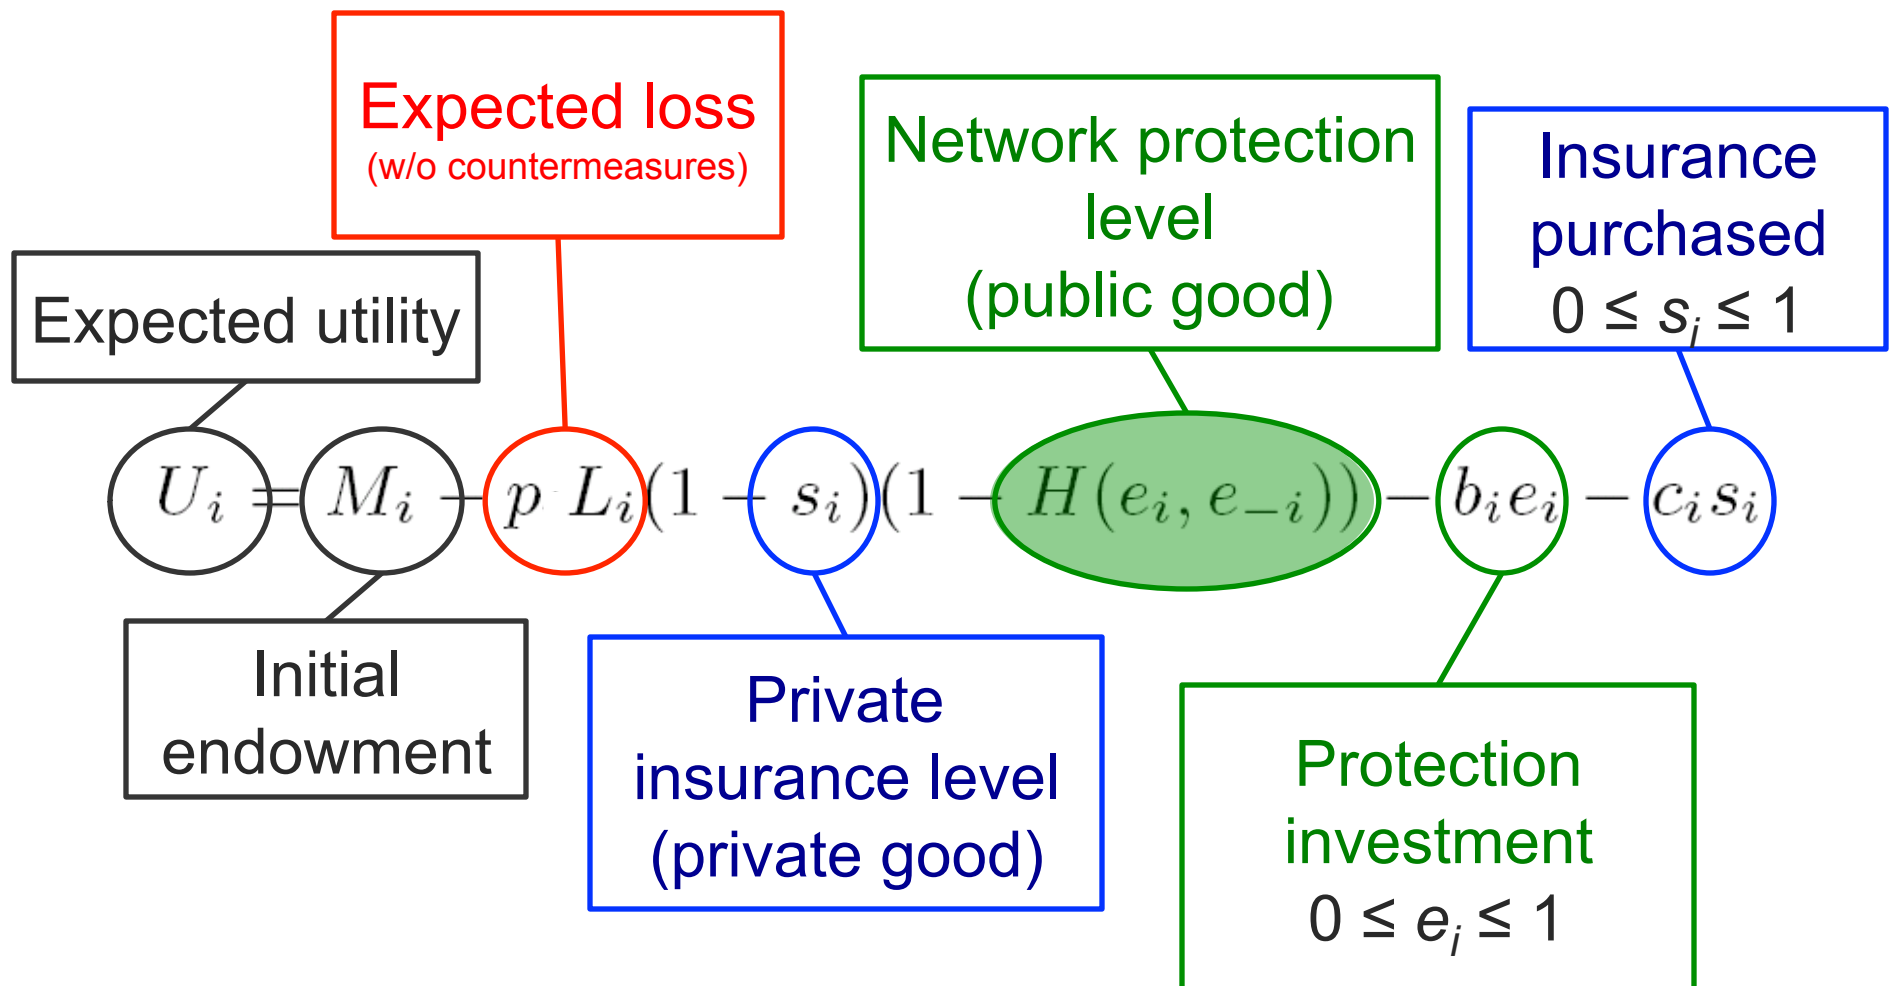
- Instead of considering security decisions to be determined by a single “security” variable
- Two key components of a security strategy
 - Self-protection (e.g., patching system vulnerabilities)
 - Joint protection level determined by all participants of a network.
 - Public good
 - Self-insurance (e.g., having good backups)
 - Individual level of loss reduction
 - Private good

General model

- N players with per-round endowment M_i
- Attacks arrive with an exogenous probability p
 - $0 \leq p \leq 1$ [uniformly distributed]
 - Loss $L_i > 0$
- Player i chooses a self-insurance level s_i and a self-protection level e_i
 - $0 \leq s_i \leq 1$ and $0 \leq e_i \leq 1$
 - e_{-i}, s_{-i} denote others' investments
- Nominal costs of self-protection b_i and self-insurance c_i
 - $b_i > 0$ and $c_i > 0$
- H is a group security contribution function of e_i, e_{-i}
 - H is assumed to be defined for all values over $(0, 1)^N$

$$U_i = M_i - p \cdot L_i (1 - s_i) (1 - H(e_i, e_{-i})) - b_i e_i - c_i s_i$$

General model



Determination of utility

- Three factors decide magnitude of a loss
 - Whether an attack takes place (p_i)
 - Whether the individual invested in self-insurance ($1-s_i$)
 - Magnitude of the joint protection level ($1-H(e_i, e_{-i})$)
- Self-insurance *always* lowers the potential loss
- Protection *probabilistically* determines whether an attack is successful
- All decisions are made simultaneously
- Equation yields an *expected* utility


$$U_i = M_i - p_i L_i (1 - s_i) (1 - H(e_i, e_{-i})) - b_i e_i - c_i s_i$$

Three traditional contribution functions

- Tightly coupled networks
- Total/average effort
- Weakest-link
- Best shot

Total effort

- Attacker needs to conquer network one-by-one
 - E.g., attacker wants to slow down distributed transfer of a file on a p2p network

$$H(e_i, e_{-i}) = \frac{1}{N} \sum_i e_i$$


Weakest-link

- Attacker only needs to find least protected node
 - Computer security: “As computer networks are cobbled together, the Law of the Weakest-Link always seems to prevail.”
- Perfect complements




$$H(e_i, e_{-i}) = \min(e_i, e_{-i})$$

$$U_i = M_i - p L_i(1 - s_i)(1 - \min(e_i, e_{-i})) - b_i e_i - c_i s_i$$

Best shot: The last stand

- Attacker needs to compromise most secure node
 - Information availability: Censorship-resistance & resilience
 - Communication availability: Destruction of all paths between two nodes in well-connected network

$$H(e_i, e_{-i}) = \max(e_i, e_{-i})$$


$$U_i = M_i - p L_i(1 - s_i)(1 - \max(e_i, e_{-i})) - be_i - cs_i$$

- Maximum effort

New models

- Loosely coupled networks
- Weakest-target security game
 - Without mitigation
 - With mitigation
- An attacker is interested in securing access to an arbitrary set of entities with the lowest possible effort
 - Select machines with the lowest security level
 - Good strategy if return is relatively low, e.g., spam distribution
 - Attacker with limited skills, e.g., if merely using automated attack toolboxes

Weakest target security game (without mitigation)

- Attacker **always** able to compromise the entity (or entities) with the lowest protection level
- Other entities stay unharmed

$$H(e_i, e_{-i}) = \begin{cases} 0 & \text{if } e_i = \min(e_i, e_{-i}) \\ 1 & \text{otherwise} \end{cases}$$

$$U_i = \begin{cases} M_i - p L_i(1 - s_i) - b_i e_i - c_i s_i & \text{if } e_i = \min(e_i, e_{-i}) \\ M_i - b_i e_i - c_i s_i & \text{otherwise} \end{cases}$$

Weakest target security game (with mitigation)

- Probability that the attack on the weakest protected player(s) is successful dependent on the security level $\min(e_i)$ chosen

$$H(e_i, e_{-i}) = \begin{cases} 1 - e_i & \text{if } e_i = \min(e_i, e_{-i}) \\ 1 & \text{otherwise} \end{cases}$$

Social optimum vs. Nash equilibrium

- **Social optimum:** set of strategies that maximizes total (network) utility $U = \sum_{i=1}^N U_i$
 - Ideal configuration for the community
 - What a benevolent government would want
- **Nash equilibrium:** set of strategies in which no individual player can increase their individual utility U_i by changing their protection or insurance settings
 - Selfish equilibrium
 - Best response to others' actions

Intuition behind Nash equilibrium outcome

- 3 types of pure Nash equilibria in our games
 - Protection only $(e_i, s_i) = (e^0, 0)$ (w/ $e^0=1$ fairly common)
 - Insurance only $(e_i, s_i) = (0, 1)$
 - Inactivity $(e_i, s_i) = (0, 0)$
- Increasing network size N affects Nash existence/nature

Total effort

- **Full protection eq.:** Protection is cheap, potential losses are high, and insurance is extremely overpriced
 - If $pL > bN$ and $c > b + pL(N-1)/N$
- **Full self-insurance eq.:** Expected losses above insurance costs, and protection is expensive
 - In other cases with $pL > bN$, or if $c < pL < bN$
- **Passivity eq.:** If both costs are too high
- Increasing N : Since players share contributions protection becomes more unlikely

Weakest link

- **Multiple protection equilibria $(e^0, 0)$:** everybody picks the same minimal security level, but no one has any incentive to lower it further down
 - If $pL > b$ and $\{(e^0 > (pL - c)/(pL - b) \text{ for } c < pL) \text{ and } (pL \geq c)\}$ and $b \leq c$
- **Full self-insurance eq.:** If the system is not initially secured well enough (by having all parties above a fixed level), players prefer to self-insure essentially
- **Passivity eq.:** If both cost are too expensive
- Increasing N : Likelihood of defection rises

Best shot

- **Full protection eq.:** No symmetric Nash equilibrium.
 - Players have strong incentive to freeride on others efforts
 - High cost of lack of coordination
- **Full self-insurance eq.:** Does exist if $b > c$ and $pL > c$.
- **Passivity eq.:** If $pL < b$ and $pL < c$

- Increasing N : Independent of network size since no protection equilibrium

Weakest target (without mitigation)

- **Pure Nash:** Equilibria for non trivial values of b , p , L and c do **not** exist
- **Mixed Nash:**
 - Probability distribution function of self-protection

$$f(e) = \frac{1}{N-1} \frac{b}{pL} \left(\frac{c-be}{pL} \right)^{\frac{1}{N-1} - 1}$$

- Likelihood of full insurance: $q = F(0) = 1 - \left(\frac{c}{pL} \right)^{\frac{1}{N-1}}$

- Increasing N : Likelihood of insurance drops; players prefer to gamble and hide in crowd

Weakest target (with mitigation)

- **Full protection eq.:** Individuals are only safe from attack with full effort
- **Mixed Nash:**
 - Does exist, with similar distribution function

$$f(e) = \frac{1}{N-1} \left(\frac{(b-c)pL}{pL^2(1-e)^2} \right) \left(\frac{c-be}{pL(1-e)} \right)^{-\frac{N-2}{N-1}}$$

- Likelihood of full insurance: $q = F(0) = 1 - \left(\frac{c}{pL} \right)^{\frac{1}{N-1}}$

Summary of homogeneous results

$(L_i = L, b_i = b, c_i = c, M_i = M, \text{ pure Nash})$

Protection

Self-Insurance

Total Effort	$pL > bN$ and $c > b + pL(N-1)/N$	Other cases with $pL > bN$ or $c < pL < bN$
Weakest Link	Multiple symmetric protection equilibria	$pL > c$ and too high protection cost
Best Shot	No symmetric Nash	Does exist if $b > c$ and $pL > c$
Weakest T w/o M	No Nash	<i>No Nash</i>
Weakest T with M	<i>Full protection if $b \leq c$</i>	<i>No Nash</i>

Role of a Social Planner (1)

- To achieve a social optimum
 - Sum of all players' utilities is maximized
 - Benevolent dictator
- Total effort:
 - More self-protection eq. ($pL > bN$)
- Weakest-link:
 - Planner would choose highest protection level
 - Pareto-optimal
- Best shot:
 - Planner now selects full protection for exactly **one** individual
 - In Nash eq. individuals frequently failed to protect
 - Insurance not needed

Role of a Social Planner (2)

- Weakest target without mitigation:
 - Sacrificial lamb
 - E.g., Honeypot
 - With or without insurance
- Weakest target with mitigation:
 - Planner only sacrifices the lamb if insurance costs are high
- Note: Contributions under social planner rule are sometimes lower



Impact of heterogeneity

- Both the homogeneous and heterogeneous cases are relevant to security analysis
 - Lack of security through diversity
 - Boat anchors
 - Security through diversity
 - Protocol randomization
 - Distribution of distinct software modules
 - Attackers overcoming diversity
 - Chameleonic threats (e.g., worms that are also email viruses)
 - Cross-platform exploits

Total effort – Decreased fragility

- More restricted player(s) have impact on dynamics of the game
 - *In 2-player game players can dissuade others from protecting*
- Each player reacts to changes in average protection over $(N-1)$ players
 - Necessary condition for domino effects to occur (from protection to self-insurance):

$$\left| \max_{2 \leq i \leq K} \frac{c_i - b_i}{p_i L_i} - \min_{2 \leq i \leq K} \frac{c_i - b_i}{p_i L_i} \right| < K - 1$$

Increased heterogeneity improves stability of protection outcome

Weakest link – Increased fragility

- Likelihood of full protection eq. is conditioned by player with largest difference between protection and insurance costs
 - Breakdown if $b_i > c_i$ for any player i
 - Protection eq. $(\hat{e}(0), 0)$ feasible if the minimum protection value chosen by any player i fulfills:

$$\hat{e}(0) > \max_{1 \leq i \leq N} \{ (p_i L_i - c_i) / (p_i L_i - b_i) \}$$

Likely to be harder and harder to meet as N grows

Best shot – Increased fragility

- Protection eq.: All players free-ride on efforts of one player i
 - If there exists a unique player i with $b_i < c_i$ and all other players choose initially at most protection levels with

$$\max\{e_{-i}(0)\} < 1 - b_i/p_i L_i$$

- Probability ρ that a protection eq. is reached given a network size N ,

$$\rho \leq N x_k^{N-1} \prod_{j \neq k} \left(1 - x_j^{N-1}\right)$$

with $x_i = F\left(1 - \frac{b_i}{p_i L_i}\right)$ (conditioned by initial security distribution, less than 1)

Declines to zero with increasing N

Weakest target (without mitigation)

- **Pure Nash:** Equilibria for non trivial values of b , p , L and c do **not** exist
- **Mixed Nash:**
 - Probability distribution function of self-protection

$$f(e) = \frac{1}{N-1} \frac{b}{pL} \left(\frac{c-be}{pL} \right)^{\frac{1}{N-1}-1}$$

- Likelihood of full insurance: $q = F(0) = 1 - \left(\frac{c}{pL} \right)^{\frac{1}{N-1}}$

- Increasing N : Likelihood of insurance drops; players prefer to gamble and hide in crowd

Weakest target (with mitigation)

- **Full protection eq.:** Individuals are only safe from attack with full effort
- **Mixed Nash:**
 - Does exist, with similar distribution function

$$f(e) = \frac{1}{N-1} \left(\frac{(b-c)pL}{pL^2(1-e)^2} \right) \left(\frac{c-be}{pL(1-e)} \right)^{-\frac{N-2}{N-1}}$$

- Likelihood of full insurance: $q = F(0) = 1 - \left(\frac{c}{pL} \right)^{\frac{1}{N-1}}$

So, what did we learn?

- Important to know which game one is playing
 - Results give you insight whether monoculture vs. heterogeneity is desirable
 - Where to invest/give incentives
- Leverage options for redesign of organization to increase security incentives under consideration of most likely security threat
- Utilize strategic uncertainty:
 - Social planner might reduce contributions
 - Consider different network sizes
- Self-insurance desirable security primitive?
 - Agents often fail to protect
 - Notable differences to Varian (2000)

Sensitivity analysis and limitations

- Simplified assumptions on costs and distribution of attacks
- Different tie-breaking rules and game modifications
 - E.g., weakest-target game with K machines compromised
 - Hirshleifer: *"The total of the best three shots, or the average of the best and worst shots, or the variance or skewness."*
 - Weaker link and better shot

Limited information

- Limited information environment
 - Agents know little about
 - Structure of game and size of network
 - Payoffs and cost effectiveness
- How does game change when limited information is taken into account?
 - Create a two-tiered network of players
 - Distinguish between security experts
 - who have limited information but understand the game
 - and naïve users
 - who have limited information and don't understand externalities

Refinements

- Complete vs incomplete information
 - An expert with complete information knows the expected losses for all players.
 - An expert with incomplete information knows her own expected loss L_i but does not know the expected losses of other players.
 - Experts assume that expected losses are independently and uniformly distributed in $[0,1]$.
- Expert vs. naïve players
 - Expert players know the contribution function H and understand its effects.
 - Naive players are myopic; they behave as if

$$H(e_1, \dots, e_n) = e_i$$

Methodology

- The question: to what extent does information security expertise help to make a network more secure?
- The methodology:
 - For each game and information condition, we derive conditions for existence of symmetric (Bayesian) Nash equilibria as a function of the protection cost b and the number of expert players k .
 - Where these equilibrium conditions are met, we compute expected utilities for all players, as well as the overall security outcome.
 - Finally, we determine the configuration yielding the expected social optimum, and we propose a system of side payments between experts that would facilitate this configuration.

Results

- In the Best Shot game, experts have a strong incentive to free-ride (Tragedy of the commons). Adding experts decreases the likelihood that the network is protected.

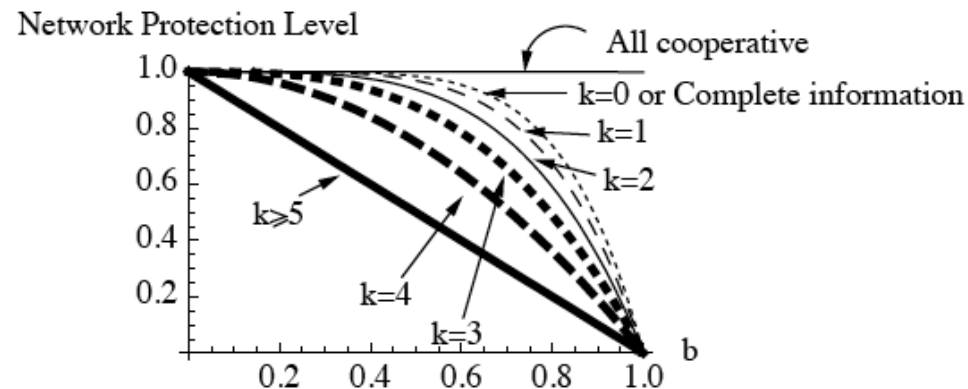


Fig. 2. Best shot. Evolution of the network protection level as a function of the protection cost b . The different plots vary the number of experts k in a network of $N = 6$ players. We observe that the fewer experts participating in the game, the higher the network protection level is, on average.

Results

- Protection equilibria in the Weakest Link game only exist when protection costs are small; and the problem is exacerbated by the addition of expert players.


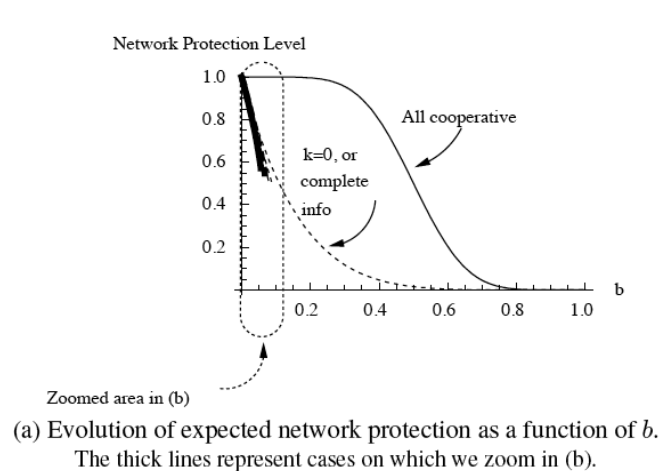


Fig. 3. Weakest link. Evolution of the network protection level as a function of the protection cost b . The short lines illustrate the presence of limiting conditions on protection equilibria for this game. Where the lines end, the expected network protection level becomes zero. Also note that the cases $k = 1$ and $k = 2$ produce identical curves.

Results

- In the Total Effort game, the individual benefit of an investment is always proportional to a $1/N$ fraction of the investment's cost, regardless of the actions of other players. Experts understand this feature and do not protect very often.

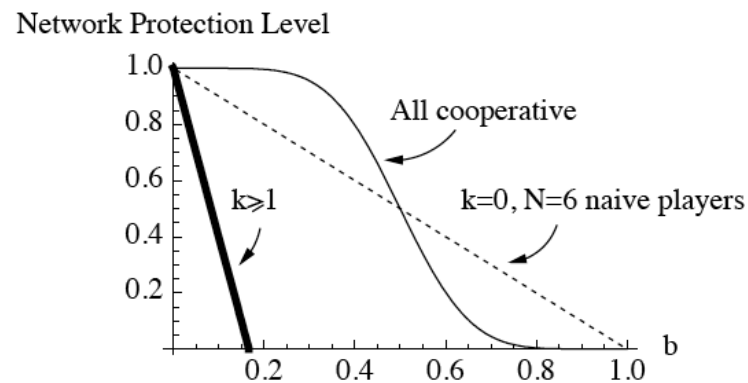


Fig. 4. Total effort. Evolution of the network protection level as a function of the protection cost b . For any number of experts $k \geq 1$, the network protection level is inferior to that obtained with a network solely consisting of naïve players. The cooperative equilibrium, here, provides a less desirable overall system outcome as soon as b exceeds 0.5.

Implications

- (In some contexts), security experts are useful when (and only when) they collaborate.
- When security is divided among independent agencies, it is important to develop mechanisms for facilitating interagency collaboration.
- User education should focus on the collaborative nature of security.

Summary

- Externalities critical in security games
- Availability of insurance in general harmful to obtaining satisfactory protection levels
 - Need to have insurance pricing dependent on protection level (similar to home or car insurance)
 - How do we measure security achieved?
- Security expertise may actually degrade the total level of security achieved in the network
 - Limited information may actually not even impact non-expert players as much, especially if they don't understand externalities
- Need to complement this work with behavioral studies
 - Users are not perfectly rational, but on the other hand, to be very vulnerable to immediate gratification – need to parameterize these insights

Thank you!

Nicolas Christin,
CMU INI & CyLab
nicolasc@cmu.edu

<http://www.andrew.cmu.edu/user/nicolasc>